

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

Continuous Bandwidth Assessment and Feedback for Voice-Over-Internet-Protocol (VOIP) Comparing Packet's Voice Duration and Arrival Rate

Background of Invention

[0001] This invention relates to voice-over-Internet-Protocol (VoIP) systems, and more particularly to measurement of current bandwidth of VoIP channels on an unregulated network such as the Internet.

[0002] The widespread availability of the Internet has allowed some traditional applications such as telephone calling to use the Internet rather than traditional telephone networks. Voice-over-Internet-Protocol (VoIP) applications capture a user's voice, digitize and compress the voice, and transmit the coded voice as data inside Internet-protocol (IP) packets. The VoIP packets can be sent over the Internet like any standard IP packet.

[0003] VoIP applications can be installed on personal computers (PC's), other devices connected to the Internet, or on translation servers such as Internet-to-Telephone gateways. Each party to a call runs a local copy or client of the VoIP application. Each VoIP application captures and sends voice data, and receives VoIP packets that are decoded and played to the local user. Thus full-duplex voice calls can be made by exchanging VoIP packets between peer-to-peer client applications.

[0004] Figure 1 is a diagram of a prior-art VoIP system experiencing packet loss. VOIP application 10 is operated by user A while VOIP application 12 is operated by user B at

different nodes on the Internet. User A's speech is digitized, coded, compressed, and fitted into IP packets 20 by VOIP application 10. These IP packets 20 containing user A's voice are routed over the Internet to VOIP application 12. VOIP application 12 receives these IP packets 20, extracts and de-compresses the voice data, and plays the voice as audio to user B. User B's voice is then captured, captured, coded, compressed, and fitted into IP packets 22 by VOIP application 12. IP packets 22 containing user B's voice are also routed over the Internet back to VOIP application 10 for playback to user A. Thus a full-duplex voice call can be made over the Internet using applications 10, 12.

[0005] IP packets can be routed over a wide variety of paths using the Internet. Indeed, the de-centralized nature of the Internet allows routing decisions to be made at a number of points along the paths between applications 10, 12. The paths taken by packets 20 in the A-to-B direction can differ from the path taken by packets 22 in the reverse (B-to-A) direction. For example, packets 20 may pass through intermediate routers 14, 16, while packets 22 pass through router 18. Such non-symmetric routing can produce non-symmetric routing delays and challenges for the VOIP system.

[0006] Various network problems may occur. A router may temporarily fail, causing some packets to be delayed or lost entirely. The number of arriving packets may suddenly jump, producing congestion such as at router 18. Router 18 may delay packets 24 while the increased packet load occurs. Packets may continue to be delayed after the initial failure is fixed as the packet backlog is worked off. If the input buffers for router 18 overflow, packets 24 may be dropped or lost rather than simply delayed.

[0007] Bandwidth limitations may also occur. Packets may need to reach a user through a low-bandwidth dial-up modem line. Occasional interference may further delay packets. The modem user may send email or browse a web site, reducing further the limited bandwidth available to the VOIP application's packets. Thus bandwidth limitations may be both permanent and temporary.

[0008] Figure 2A shows voice data that is packetized and transmitted. The user's voice can be captured as analog waves of varying frequencies that are digitized and coded. The coded voice data is divided into packets and transmitted. Sequence numbers are added to the packets to allow the packets to be re-ordered when some are delayed

more than others. The sequence numbers thus allow for out-of-order reception. In this example the coded voice is divided into four packets, each packet containing coded voice data for an equal, fixed time period of 20 milli-seconds.

[0009] Figure 2B shows packetized voice data received after varying network delays. The sequence numbers are used to re-order the packets when they arrive with varying network delays.

[0010] In this example, packet 2 is delayed slightly, causing a gap to occur between the end of playing the voice for packet 1, and the start of voice play for packet 2. A larger gap occurs between packets 2 and 3, between times S2 and S2'. These gaps may be filled in by interpolating voice data, or by adding silence. However, the pace of the user's voice may seem uneven or jerky due to such gaps.

[0011] Of course, all voice could be delayed by a large amount, such as 5 seconds, to allow for late packets. However, this requires a larger packet-input buffer and would greatly increase the delay or latency that the user hears. This delay may be noticeable to the user and annoying. Full-duplex conversation becomes impractical as the delay grows to several seconds. Thus the input buffer has a practical size limit, and packets cannot be delayed for too long.

[0012] Such gaps caused by delayed packets can reduce the quality of the voice played. When a temporary interruption occurs along the path taken by the VoIP packets, packets may pile up in buffers near the point of interruption. Should service be quickly restored, the stored packets in the buffers may be sent after some delay. However, longer-duration interruptions can cause router buffers to overflow. Packets may then be dropped or discarded before reaching their destinations.

[0013] Once the interruption ends, the older packets are likely to be sent first by the router. Newer packets may be delayed even after the interruption ends as the backlog of packets is transmitted. Thus stale packets of older voice data may be delivered before more current voice data. These older packets may already be too old to be played, resulting in a lengthening of what was a brief moment of congestion.

[0014] Detecting when such congestion occurs or when a limited bandwidth is available could be useful. Transmission of voice packets could be paused to prevent

[0016] During initialization of a call between applications 10, 12, some provisioning may be performed to determine the initial bandwidths available between applications 10, 12. Such provisioning may be similar to fax machines that negotiate compression standards used and bandwidth or baud rate for each call. However, changes to the Internet that later occur during the call are not detected once provisioning is over and the call is started.

Brief Description of Drawings

[0021] Figure 3 is a diagram of a VOIP system that continuously measures incoming-packet bandwidth and transmits bandwidth estimates in outgoing packets.

[0023] Figure 5 shows an outgoing VOIP packet with bandwidth and congestion estimates

analyzed by bandwidth detector 40 to determine the packets' travel time along path 39 and indirectly estimate the bandwidth of path 39. This bandwidth estimate from bandwidth detector 40 is added to outgoing packets 34. Packets 34 contain both voice data from user A, VA, and the bandwidth estimate for packets 36 sent by user B, BW_B.

[0032] When packets 34 are received by VOIP application 32, user A's voice data VA is extracted and played back to user B, and the bandwidth estimate BW_B is read, allowing VOIP application 32 to adjust or halt its transmission of outgoing packets 36. For example, when bandwidth is reduced, VOIP application 32 can signal user B of the problem, such as by generating an audible beep to indicate the poor bandwidth.

[0033] Bandwidth detector 42 in VOIP application 32 also measures the arrival rate of incoming packets 34 to estimate the bandwidth of path 38. This bandwidth estimate for user A, BW_A, is added to outgoing packets 36 which contain user B's voice data, VB. Thus packets 36 contain VB and BW_A, while packets 34 contain VA and BW_B.

[0034] Bandwidth detector 42 in VOIP application 32 also measures the travel time or latency of incoming packets 34 to estimate the congestion of path 38. When latency begins to increase, congestion is starting to appear.

[0035] One-Way Latency Measured, Not Round-Trip Time

[0036] The latency or travel time measured by bandwidth detector 40 is not the round-trip travel time. The round-trip travel time includes both paths 38, 39. Instead, only the one-way latency is measured, from VOIP application 32 to VOIP application 30 over path 39. Separate bandwidth and congestion estimates allow for asymmetric latencies, such as when path 38 is restricted while path 39 is not. More precise bandwidth estimates are thus possible.

[0037] Figure 4 shows in more detail a VOIP application with a bandwidth detector. VOIP application 32 captures user B's voice and stores the digitized voice as voice data 54. Codecs 52 are one or more voice encoders that compress and encode the raw digitized voice using a variety of algorithms. Standard as well as proprietary codecs can be used. Packetizer 50 forms the outgoing IP packets by adding headers and catalogs of the voice data, to the encoded voice data from codecs 52.

[0038] Incoming packets with user A's voice data are received and stored by jitter buffer 48. Some delay and variation in packet reception is accommodated by jitter buffer 48, and packets can be re-ordered by sequence number if received out of order. The packets are sent to core manager 56 of VOIP application 32, which extracts the voice data from the packets, examines the voice catalog, and selects the specified codec to decode and decompress the voice data. The final decoded, decompressed voice data is played as audio to user B. Core manager 56 may contain a variety of software modules including a user interface or may call other modules, library, or operating system routines.

[0039] Latency Measured by Time-Stamps

[0040] Time stamper 46 provides time-stamps or clock values that are an indication of time. Time stamper 46 generates the arrival time for each packet received by jitter buffer 48. Each packet also contains a send time that was included by the other VOIP application. Bandwidth detector 42 compares the arrival time with the send time for each packet to get the packet's travel time or latency. The change in latency over time is used to determine when congestion occurs.

[0041] The arrival rate of incoming packets is used to estimate bandwidth. For example, when the arrival times between packets increase, bandwidth is reduced. Bandwidth detector 42 generates current estimates for the incoming bandwidth, BW-EST, and congestion, CONG-EST.

[0042] Packetizer 50 receives the bandwidth and congestion estimates from bandwidth detector 42 and adds these to outgoing packets. The estimates may be numerical values such as 5-bit or 8-bit binary numbers that represent a magnitude of bandwidth or congestion, or may be more qualitative values such as 2 or 3-bit values that indicate "good", "average", "poor", or "blocked" paths. One-bit values such as a congestion flag may also be used.

[0043] When packets fail to arrive at jitter buffer 48, or are substantially late, such as more than 2 seconds, the packet loss counter is incremented. The packet loss counter PKT-LOSS may also be included in outgoing packets.

[0044] Figure 5 shows an outgoing VOIP packet with bandwidth and congestion estimates

for the incoming path. IP packet 36 includes network-level header information such as a Telnet-Connect-Protocol (TCP) or user datagram protocol (UDP) header. Ethernet and Internet Protocol (IP) information may also be included. IP packet 36 may be further encapsulated during routing, such as by adding Virtual Private Networking (VPN) or other transport layering. Headers for lower-layer protocols can encapsulate headers for higher-level protocols.

[0045] IP header 60 contains the destination and source IP addresses while TCP/UDP header 62 contains the TCP or UDP port or other TCP information. Checksums and other information may also be included. Application audio or voice data field 68 contains the compressed and encoded voice data and may be sub-divided into several sub-fields.

[0046] Send time field 64 contains the send time $S(N)$ or time-stamp value placed into packet 36 when the packet was transmitted. Catalog 66 is a directory of the voice-data contents of voice-data field 68. The playing time for the voice data, such as 20 milli-seconds, is the duration $D(N)$. This voice duration can be explicitly or implicitly contained in catalog 66. The duration may have to be calculated by adding durations of segments of voice data in voice-data field 68, or by considering the kind of codec and compression used and the number of bytes of voice data.

[0047] The bandwidth estimate from the bandwidth detector can be added to packet 36. For example, the bandwidth estimate BW-EST, congestion estimate CONG-EST, and packet-loss counter PKT-LOSS can be added to the end of packet 36. Often unused bits are available at the end of the compressed voice data in voice-data field 68, or additional bits can be added to packet 36 for estimate fields 70, 72, 74, which contain the bandwidth, packet-loss, and congestion values.

[0048] Figure 6 highlights time-stamping of arriving packets. VOIP packets 76, 77, 78 arrive from the Internet and are stored in jitter buffer 48. Each packet N contains a send time $S(N)$ and a voice duration $D(N)$. The voice duration may be explicit or implicit. For example, the total voice duration may have to be calculated as the sum of the durations of data sub-fields or audio frames in a packet, or may have to be adjusted for different codings and codecs used.

[0049] As packets 76, 77, 78 arrive, time stamper 46 outputs a value for the current time, which is associated with each arriving packet. For example, packet 76 arrives or is received by jitter buffer 48 at time $R(1)$, while packet 3 is received at time $R(3)$. These reception-time values can be stored with the packets in jitter buffer 48, or may be stored in a separate memory or buffer area but be associated or linked to the packet. The send time and duration from each packet could also be extracted and stored with the reception time in a different memory, such as one accessed by the bandwidth detector.

[0050] Congestion Detected by Latency Changes

[0051] The one-way latency or travel time is the difference of the send and reception times. Packet N 's latency is $R(N) - S(N)$. For actual networks, the latencies vary. When latency increases, congestion may be occurring. When latencies drop, congestion may be easing. The packet's latency is compared to a moving average of the latencies of many packets to determine when latency is increasing or decreasing, and thus signal when congestion is increasing or decreasing.

[0052] Bandwidth measured by Arrival Rate and Voice Duration While latency changes are used to signal congestion, packet arrival rates are used to determine bandwidth. A packet's voice duration should equal the time between packet arrivals. Under ideal network conditions, the time between successive packets is equal to the voice duration. For example, when packets contain 10 milli-seconds of voice, the packets need to be sent every 10 milli-seconds (ms) for a continuous voice transmission. If packets contains 50 ms of voice, then it is expected to arrive 50 ms after the previous packet.

[0053] The time between arrivals of packets with successive sequence numbers is the inter-packet arrival time. This inter-packet arrival time is compared to the voice duration of the most recent packet to arrive. When the inter-packet arrival time is greater than the packet's voice duration, the network is too slow. When a network recovers or speeds up, inter-packet arrival times can be less than the packets' voice durations.

[0054] Figures 7A-C are flowcharts highlighting estimating bandwidth and congestion

from packet arrival rates, latencies, and voice durations. In Fig. 7A, when a packet arrives at the jitter buffer, the jitter buffer or associated logic reads the packet's sequence number and determines if the packet is excessively late, such as more than 2 seconds late, step 102. When the packet does not arrive within the time limit, the packet loss counter PKT-LOSS is incremented, step 104. Packets that never arrive, such as packets that are dropped by the network, also increment the loss counter.

[0055] When a packet arrives within the time limit, step 102, bandwidth estimation 100 is performed as shown in Fig. 7B. Congestion estimation 120 as shown in Fig. 7C is also performed. These estimations can be performed on each arriving packet as the packet arrives or soon after, or can wait until several packets have arrived and can be processed together, or can be processed periodically at a set time interval or in the background when processing time is available.

[0056] Each packet's reception time $R(N)$ is generated by the time stamper, and the packet's send time $S(N)$ is extracted from the packet. Each packet's voice duration $D(N)$ is also determined. In fig. 7B, bandwidth estimation 100 determines the inter-packet arrival time DT , which for packet N is $R(N) - R(N-1)$. Packet $N-1$ can be the packet with the previous sequence number before packet N . Once the inter-packet arrival time DT is calculated, step 106, it is compared to packet N 's voice duration, $D(N)$.

[0057] When the inter-packet arrival time DT is less than the voice duration $D(N)$, the packet arrived early, step 108. This indicates that the network is operating more efficiently than currently estimated, and may be recovering from an earlier network problem or constriction. Since the current bandwidth estimate underestimates the potential bandwidth, the bandwidth estimate $BW-EST$ is increased, step 110. While the bandwidth estimate could be increased by a fixed amount or some other amount, in this example $BW-EST$ is increased in proportion to the absolute value of the fraction $(R(N) - R(N-1) - D(N)) / D(N)$, which is also $(DT - D(N)) / D(N)$, or the excess of the inter-packet arrival time DT over the voice duration, divided by the voice duration. The $BW-EST$ may be increased by the whole fraction, or by a portion such as 10% or 50%. The portion may be programmably changed or dynamically changed in some embodiments.

[0058] When the inter-packet arrival time DT is greater than the voice duration $D(N)$, the

packet arrived late, steps 108, 112. This indicates that the network is operating less efficiently than currently estimated, and may be suffering from a network problem or bandwidth constriction. This can occur on limited-bandwidth links such as a modem line when the user sends or receives email or browses a web site while also using the VOIP application.

[0059] Since the current bandwidth estimate over-estimates the true bandwidth, the bandwidth estimate BW-EST is reduced, step 114. While the bandwidth estimate could be decreased by a fixed amount or some other amount, in this example BW-EST is decreased in proportion to the absolute value of the fraction $(R(N) - R(N-1) - D(N)) / D(N)$, which is also $(D(N) - DT) / D(N)$, or the excess of the voice duration over the inter-packet arrival time DT, divided by the voice duration.

[0060] When the inter-packet arrival time DT is equal to the voice duration D(N), the packet arrived on time, step 112. This indicates that the network is stable and operating as efficiently as the current estimate. The bandwidth estimate is increased by a small amount, step 116, such as 0.1 %. Increasing the bandwidth estimate when the network is stable allows the VOIP application to test if additional bandwidth is available.

[0061] In Fig. 7C, congestion estimate 120 is performed. The packet's latency or travel time from the remote VOIP application to the local VOIP application is determined, such as the difference of send and receive times, $R(N) - S(N)$. A moving average of the packet latency is kept, such as for the last 20 or 100 or 1000 packets. The current packet's latency can be added to the moving average and the oldest moving average dropped either before or after comparison.

[0062] The current packet's latency is compared to the latency moving average, step 122. When the current packet's latency is below the moving average, step 124, then the latencies are falling and the network is improving. Latencies often fall when the network is recovering from a delay caused by congestion at a routing point. Since the network is likely recovering from a problem, the congestion estimate CONG-EST is left unchanged, step 128. This allows more time for the network to stabilize.

[0063] When the current packet's latency is above the moving average, steps 124, 126,

then the latencies are rising and the network is deteriorating. Latencies often rise quickly when the network is just starting to see delays caused by congestion at a routing point. The congestion estimate is increased by a portion of the amount that the current packet's latency is above the moving average, step 130. The congestion estimate can quickly detect network problems such as at the very start of congestion using this method.

[0064] When the current packet's latency is about equal to the moving average latency, step 126, the network is stable and congestion is not apparent. The congestion estimate can be reduced by a small amount, step 132, such as 0.3 % or 0.1% or a larger value such as 1%. This allows the congestion estimates to drop back after congestion ends once the network stabilizes again. Since many packets can arrive in a short time, the congestion estimate can recover quickly even when a small change is made.

[0065] The next packet arrival can then be processed by setting packet N+1 to be packet N, and the process repeated from Fig. 7A.

[0066] Figures 8A–B show graphs of packet arrivals and bandwidth estimates. Fig. 8A has the voice time or packet sequence number as the y-axis and the actual arrival times of packets as the x-axis. In this example packets have the same voice durations and should all arrive with the same inter-packet arrival time and thus fall along ideal line 250.

[0067] During time period 200, packets arrive along ideal line 250. Fig. 8B shows that the bandwidth estimate is increased slightly during this time of network stability. However, at time period 202, packets are delayed and arrive with longer inter-packet arrival times. Arrival times T4 and T5 are delayed, causing packets to arrive below ideal line 250, with a lower slope or arrival rate.

[0068] The bandwidth estimate is reduced by a portion of the lateness, and falls sharply during time period 202. When packets are very late, the bandwidth estimate can be reduced even before the packet arrives. A timer can wake up periodically to examine the most-recently-arrived packet. The maximum-size packet's duration can be compared against the time that has transpired since the last packet arrival. In an

example where the network comes to almost a complete halt for an extended period, late packets can be detected by expiration of a maximum inter-arrival time. This can be factored into the bandwidth and congestion estimates.

[0069] Packets begin arriving at the ideal rate during period 204. The packets have the same slope as ideal line 250, but are below line 250 due to the delays from period 202. The bandwidth estimate rises slightly during this period.

[0070] The network recovers quickly during period 206 as many packets arrive in a short time. This can occur as a router recovers from a delay and works off its packet backlog. The packets rapid arrival produces a slope higher than that for ideal line 250, and eventually the packets reach line 250. The bandwidth estimate rises quickly during period 206 as a portion of the difference of inter-packet arrival time and the voice duration of the voice data inside the packets.

[0071] Finally in period 208 the network is again stable and packets arrive along ideal line 250. The bandwidth estimate is edged up slightly to test the upper limit of bandwidth.

[0072] Figures 9A-B show graphs of packet latencies and congestion estimates. In Fig. 9A, latencies of arriving VOIP packets are plotted as a function of voice time. A similar graph can be made using time or sequence number for the x-axis. The dotted line is the moving average of the latencies and shows less movement than the current packet latencies since it is an average.

[0073] Latencies are rising slightly over long time periods, as shown by the upward bias to the moving average during periods 210, 214. The congestion estimate remains relatively flat during periods 210, 214.

[0074] During period 212, a network problem or constriction occurs, causing the current packet latencies to rise sharply above the moving average. This can occur when a user sends or receives email over a modem line that is being used by the VOIP packets. The congestion estimate quickly rises as the latencies rise.

[0075] Rather than fall back as quickly as the latencies as the peak ends, the congestion estimate remains high as the current latencies fall sharply as Fig. 9B shows. This flat

[0082] Several other embodiments are contemplated by the inventor. For example various

combinations of software, hardware, or firmware implementations are possible and various routines can be called and executed sequentially or in parallel. While the VOIP packets have been described as being routed over the public Internet, packets may be routed over other networks or combinations of networks such as Ethernets, Intranets, wireless networks, satellite links, etc. The audio packets can also include multi-media data such as images or text.

[0083] Rather than estimate bandwidth by calculating the latency for each packet, only a subset of the packets could be checked, such as every 5th packet or every 50th packet. The durations of intervening packets could be summed. The bandwidth and congestion estimates could likewise be embedded in only some of the outgoing packets rather than all outgoing packets. The bandwidth and congestion estimates could also be sent in separate packets without voice data. The voice data is really audio data that is often voice, but could include other audio data such as songs, music, traffic noise, etc.

[0084] The bandwidth estimate could also be kept constant when the network is stable, or could be increased by a different amount or by a variable amount. The congestion estimate could be performed before or after the bandwidth estimate, or at the same time. Parallel processing could be used on some systems.

[0085] Network recovery typically is very quick, and the congestion estimate can be raised immediately, or as shown in the previous embodiment, the congestion estimate can be left at its present level until such time as the network has cleared any backlog of stale or delayed packets.

[0086] The bandwidth and congestion estimate routines could be activated by the jitter buffer when packets are late in arriving but before the packets arrive. Since the sending times of the missing packets are not known, they may be interpolated from other packets, or a fixed number used to calculate the new arrival time, latency, or voice duration. The amount of voice data in packets can vary from packet to packet rather than be the same for all packets as described in the simplified examples. The jitter buffer may perform other functions, such as detecting and processing duplicate and missing packets. The jitter buffer can also vary the amount of buffering and consumption rate of voice data in concert with occurrences of congestion to minimize

[0090] Data values can be considered "equal" if within a certain range of each other, such as within 1% or 5% or 0.1%. Also, rounding of values can be performed before comparison, effectively providing a range of "equal" values. Congestion and bandwidth estimates can use only a few bits to indicate qualitative measurements such as "normal", "minor restriction", "major restriction", "blocked", or may use more bits to represent a quantitative estimate such as a percentage or data rate. One or both users could be notified of problems by a tone or a display message, or the estimates could be logged to a file for debugging. The application may visually display a network-quality meter to the user. The estimates fed back to the sending VOIP application could allow the sender to stop or reduce packet transmission when problems occur, or could adjust compression or coding to reduce bandwidth to match the estimate.

[0091] VOIP calls may be between two users on personal computers, or may consist of one user on a personal computer talking to a computer server or gateway which converts the call from VOIP to telephone or PBX or private IP phone system formats. The call could also be between two telephone or private IP-phone users with a VOIP segment somewhere in the middle carrying the call from one location to another over the Internet or similar unmanaged network but terminating the call at each end on a telephone or PBX or IP phone. Calls could also involve a conversation between one user on a PC or telephone or IP phone, and at the other end an automated voice response system such as a banking application, voicemail, auto attendant, talking yellow pages or other automated voice service. More than two parties may exist in multi-way calling. The VOIP application could carry one user's audio signal to and from a central conference server hosting a number of other callers.

[0092] The abstract of the disclosure is provided to comply with the rules requiring an abstract, which will allow a searcher to quickly ascertain the subject matter of the technical disclosure of any patent issued from this disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. 37 C.F.R. § 1.72(b). Any advantages and benefits described may not apply to all embodiments of the invention. When the word "means" is recited in a claim element, Applicant intends for the claim element to fall under 35 USC § 112, paragraph 6. Often a label of one or more words precedes the word "means". The

word or words preceding the word "means" is a label intended to ease referencing of claims elements and is not intended to convey a structural limitation. Such means-plus-function claims are intended to cover not only the structures described herein for performing the function and their structural equivalents, but also equivalent structures. For example, although a nail and a screw have different structures, they are equivalent structures since they both perform the function of fastening. Claims that do not use the word means are not intended to fall under 35 USC § 112, paragraph 6. Signals are typically electronic signals, but may be optical signals such as can be carried over a fiber optic line.

[0093] The foregoing description of the embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto.